

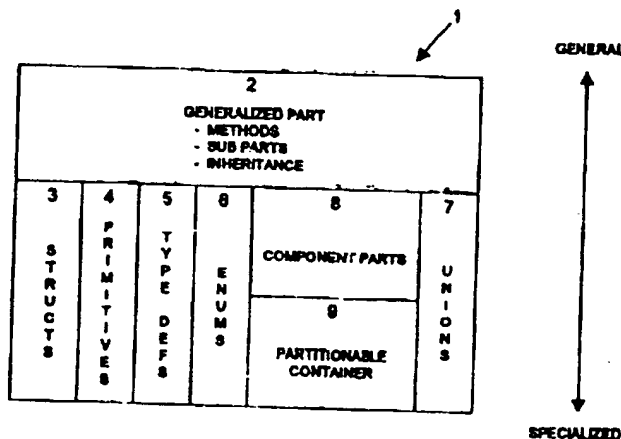


- (72) BREALEY, Chris L., CA
(72) JOHNSTON, Jeff G., CA
(72) KLICNIK, Vladimir, CA
(72) LAUZON, David M., CA
(72) LOI, Lok T., CA
(72) SEELEMANN, Alek D. II, CA
(71) IBM CANADA LIMITED - IBM CANADA LIMITEE, CA

(51) Int.Cl.⁶ G06F 9/45

(54) **MEMOIRE HIERARCHIQUE DE METADONNEES DESTINEE A
UN ENVIRONNEMENT DE DEVELOPPEMENT INTEGRE**

(54) **HIERARCHICAL METADATA STORE FOR AN INTEGRATED
DEVELOPMENT ENVIRONMENT**



(57) Dépôt de métadonnées pour un environnement de développement intégré. Le dépôt est divisé en couches de façon à définir des niveaux de comportement commun utiles à différents types d'outils de développement d'applications. Les outils d'utilisation la plus générale ont accès aux métadonnées au niveau de types à structure simple; les outils plus spécialisés ont accès à des composantes qui contiennent les propriétés d'un langage cible; les outils de haute spécialisation ont accès à des métadonnées composées à parties segmentables qui peuvent servir à construire des applications réparties.

(57) A metadata repository for use in an integrated development environment is provided. The metadata repository is layered to define levels of common behaviour useful to different types of application development tools. The most general use tools have access to metadata at the level of simple constructed types; more specialized tools have access to components that contain properties of a target language; highly specialized tools have access to composed partitionable part metadata that can be used for constructing distributed applications.

2201278

CA9-97-003

**HIERARCHICAL METADATA STORE FOR AN INTEGRATED DEVELOPMENT
ENVIRONMENT**

Abstract

5 A metadata repository for use in an integrated development environment is provided. The metadata repository is layered to define levels of common behaviour useful to different types of application development tools. The most general use tools have access to metadata at the level of simple constructed types; more specialised tools have access to components that contain properties of a target language; highly specialised tools have access to composed partitionable part metadata that can be used for constructing distributed applications.

HIERARCHICAL METADATA STORE FOR AN INTEGRATED DEVELOPMENT ENVIRONMENT

Field of the Invention

- 5 The present invention relates in general to the data processing field, and more specifically to the provision of an integrated tool development environment.

Background of the Invention

- 10 Traditional tool development follows the simple steps of creating the text-based source code, and then compiling and executing the code on a specific platform. The increasing complexity of function and integration of tools over multiple platforms means that tool development cannot generally be performed by small groups of human developers. Often the only way to produce a fully functioning application is to make use of specially-designed application development tools in order to build other tools. Today's applications are built using a mix of low level tools such as source editors and
15 debuggers, along with high level builders of user interface, data access, distribution support, and code generators. The applications built in such environments are targeted to run in a multiplicity of execution environments, on various hardware platforms, supporting many distribution mechanisms and data access mechanisms. Developing such applications in the current "state of the art" requires the use of multiple different tools for multiple different vendors, each solving a piece of the over all
20 puzzle.

- The ideal is to provide an integrated development environment (IDE) in which information provided by each tool during program development can be shared with other tools, particularly those being used in the same development effort, to avoid duplication and inconsistent interface and function
25 operations. Currently available development tools address only a small subset of this requirement. They typically focus on a source-level, single targeted application, and use a simple file system folder/directory model for their information.

CA9-97-003

Where the application development tools are tightly integrated, they can often use cross reference tables to share information about the program under development. However in an incremental development environment, the tables can be out of synch with the source code, and frequently contain less information than implied by the source code.

5

Also, any tool not closely integrated (and this is often the case with tools from different manufacturers) will have to parse the source or the tables, provided the programming language is not a barrier. This parsing exposes the tool either to parsing differences or internal table definitions.

10 Summary of the Invention

One of the key problems that needs to be addressed during construction of an IDE for a complex application is the ability of the various IDE tools and components to share information in a meaningful way. An "information store" is required to facilitate this sharing. Such a store must support usage and access characteristics of a multi tool/multi user environment and provide a flexible basis for capturing required semantic information.

15

The store must also define an overall information structure about applications under development, and organize it in a way that offers the degrees of freedom required by today's applications. These degrees of freedom include choices of targets, languages, component models, distribution, etc. The IDE further needs to define a model for sharing common application semantics among the cooperating tools, but at the same time allow individual tools to extend the shared information with private tool data. The model also need to "manage" programming "element" naming issues (naming scopes) across the various domains implied by the degrees of freedom.

20

Therefore, an object of the present invention is to provide a mechanism for organizing metadata in support of the development of complex enterprise applications by defining a layered data model for handling application development metadata, addressing the various application degrees of freedom within a single metadata store.

25

CA9-97-003

In accordance with these objects, the present invention provides a metadata container for common access tool data in an object oriented programming environment. The metadata container has a hierarchical structure that consists of simple constructed types containing subparts and encapsulated behaviour, components containing properties of a target language, and composed parts permitting partitioning for distribution.

The present invention also provides a metadata store for common access tool data in a object oriented programming environment consisting of a single base class defining common behaviour for elements in the tool data, and separate abstract class hierarchies, inheriting from the single base class, to define name scope and containment for tool data.

Brief Description of the Drawings

Figure 1 is a schematic drawing illustrating categories of metadata layered "parts" according to the preferred embodiment of the invention, and

Figure 2 is a tree diagram illustrating the model layer hierarchy according to the preferred embodiment of the invention; and

Figure 3 is a class diagram, expressed in C++, of an IDE element base class, according to the preferred embodiment of the invention

Detailed Description of the Preferred Embodiment

The preferred embodiment has been implemented as an open data model that provides source code or the common intermediate form accessible and extensible by all development tools that wish to work with a program. The data model is language independent which allows cross system source sharing. This model could be use to drive a code generator.

In the preferred embodiment, the data model provides a metadata store in support of a tool

CA9-97-003

integration framework which is fully described in our concurrently filed application titled "An Object Oriented Framework Mechanism Providing Common Object Relationship and Context Management For Multiple Tools" (IBM Docket No. CA997-002).

5 In order to handle the requirements of the range of applications, from basic source programming to component based visually constructed applications, the data model of the present invention provides a shared metadata store in which the development information is layered. Each layer defines the common behaviour required at that level. Tools using the model start with a behaviour at a given layer and extend it as appropriate for the particular tool.

10

The data model captures information about various "elements" required during the development of an application. This includes information about the logical pieces of an overall application, such as folders, files and application parts. At its highest level, the data model generally defines the common behaviour required to be recognized as an element within the integrated development environment.

15 This includes the ability for rendering the element, exposing its content (if appropriate), as well as triggering any of its actions.

This is roughly equivalent to other currently available application development information stores that map to a file system (source files and directories/folders). However the data model of the invention is further extended into key areas:

20

1. recognition of namespaces as distinct from folder containment; and
2. refinement of the application "part" concept into distinct part layers.

25 Rather than leaving it to the developer to recognize project folder containment and application part naming at the source or force the naming relationships to be the same as the folder containment, the data model of the present invention defines name scope as a "parallel" set of relationships to that of containment. This allows, for example, a part containing a C++ class A::B::C to be properly defined as a series of nested naming scopes, while at the same time allowing the same part to be used

CA9-07-003

(through containment or linking) within several project folders.

5 The data model of the invention supports application definition and assembly from parts. According to the invention, a part is a self contained piece of an overall application that is distinct from the source files that contain the source code for the part. Depending on the programming environment, a part may in fact reference several files (for example in C++, the .hpp and .cpp files) but a part is not a file container. Within the data model of the invention, a part is a separately editable piece of application definition (metadata) distinct from any source files that may be associated with the part or generated from the part definition. As will be discussed in more detail below, within the data model of the invention, all parts are named and name scoped. A part (based on its implementation) 10 may itself be a name scope for other parts, as in, for example, C++ nested classes.

The layering approach to a metadata store is schematically illustrated in Figure 1. The semantics of a part 1 are layered in a spectrum in the data model from parts used by general use tools to parts 15 used by increasingly specialised tools

At the general end of the spectrum, Generalised Parts 2 are base parts corresponding to the semantics of a constructed type within the target language. The list of generalised parts include subparts and encapsulated behaviour such as methods and inheritance. Typically, these correspond to a class 20 definition for object oriented languages such as C++ and Java™, or to some form of a structure, such as in COBOL.

A degenerate case of the generalised parts are various types supported by the data model. These types are further specialisations of generalised parts that capture or encapsulate the following concepts: fundamental types (eg. int, char) are represented by Primitive Types 4; records are represented by Structs 3; overlaid memory structures are represented by Unions 7; named cardinal values are represented by Enumerations (Enums) 6; and aliased parts are represented by Type Definitions (Type Def) 5. 25

CA9-97-003

Component parts 8 are base name parts with the additional support of a component model. A component part, regardless of its ultimate source language implementation, defines its external behaviour as a set of attributes which are called "properties", events raised by the part and actions which are externally triggerable behaviour. The data model captures this information as a component part and allows the generation of the appropriate target component model (eg Java Beans, IBM® VBE model, etc.). Within an IDE, this typically represents the level of support used by higher order "builder" tools. Component parts are referred to within the data model as "builder consumable" parts, in relation to Figure 2.

Partitionable containers 9 are component parts that expose part of their internal behaviour to a level that allows repartitioning of the part for distribution. Through the data model, these parts expose interactions among their aggregated subparts, as well as interactions across parts. This level of information is used to capture the application distribution structure regardless of the target distribution mechanism.

Depending on its specific requirements, a tool selects the appropriate level of parts support to store its tool part metadata within the data model. Shared behaviour is inherited from the data model, and the tool adds its own extensions as appropriate. Typically, these extensions are opaque tool-specific properties. For example, this structure allows a data access part which is a component part containing information about access to a data base, to be consumed by a user-interface (UI) builder producing a UI part. The UI builder is a composed part containing the data access part and UI controls. This can subsequently be partitioned into a client part (UI and server call) and server part (server logic) and access to data base. This is all done by editing the data model information. The same data model information can be used to generate a Java bean with RMI call to server and JDBC access to data base, or a C++ client with DCE access to a server with a local data base.

The preferred embodiment is implemented in an object oriented programming hierarchy which facilitates the sharing of parts through inheritance. Figure 2 illustrates an approach to data model

CA9-97-003

layers according to the preferred embodiment providing a class definition at each layer and illustrating how the class of metadata, at that level in the hierarchy, can be used by tools.

5 The most abstract base class, the IDE Element 20, is provided for the implementation of an IDE tree view. Using mostly pure virtual methods, it permits the tree view to display the contents of the model, launch actions against its elements (nodes), and attempts structural changes against the hierarchy without regard to the content details. Success of these operations (especially structural changes) depends on the supported semantics of the individual tree element derivations.

10 The common behaviour of rendering of tree nodes is defined so that each element within the tree view will show its name, icon, indication whether element is owned or linked, and an indication whether the element can be further expanded.

15 The basis for expanding branches of the tree view is the ability to list the contents of any one IDE element as its contained element. Thus, the IDE Element base class 20 provides a common behaviour to expose content. This is illustrated, in the preferred embodiment, in the sample C++ header file for an IDE Element base class set out in Figure 3. Each derived IDE element determines what constitutes its content. Unlike most other relationships supported within the model, the "contains" relationships have a requirement for supporting direct access to its element by name, as well as
20 providing cursor based access that maintains ordering. This is required because the tree view is expected to allow the nodes within the tree to be reordered using drag and drop operations or the selection of reordering actions. The relative element order is maintained within the model. A specific example of this is a visual builder pallet containing pallet categories, each of which contains parts available to the builder. Frequently used items would be expected to be placed at the top of the list
25 to reduce the need for scrolling.

The model also provides a set of methods at this level of abstraction that allow modifications to the tree structure without regard to element details. The success of these operations is determined by

CA9-97-003

the derived implementations. There are two mechanisms for constructing container relationships:

1. Ownership relationships indicating aggregation of the sub elements within the subject parts. Owned elements are destructed if orphaned by a parent; and
2. Reference relationships indicating soft links or shadow links to a sub element owned or aggregated elsewhere within the tree.

In all cases, a removal action on one element propagates the removal action across ownership relationship but not reference relationships. That is owned elements are destructed but in the case of linked elements only the link is removed. It should be noted that the data model does not enforce a single owning parent rule, although a great majority of owning relationships are expected to be of this type. In an owning relationship, the child is destructed when any parent is destructed or when the child is explicitly orphaned by any owning parent.

Since derived implementation may reject attempt to change the contents of the tree nodes, methods are provided to check if the desired operation is allowed. Even if the operation is allowed it may still fail for other reasons. The decision to accept or reject a content action is up to the object that is the recipient of the content, that is the parent object would be asked for permission not the child object.

In the preferred embodiment, the IDE Element base class provides additional common behaviour, including:

1. action support: the model provides a mechanism for triggering edit actions against the model objects. The term "edit" simply means any tool specific action. Each tool is expected to provide an implementation of the edit methods in their concrete class derivations

2. code generation support: each element can cause source output to be generated for the model metadata. A default implementation of generate () is to simply return. If output must be produced, the derived implementation need to over ride the generate () method. This generate () is a context sensitive method.

3. "dirty" part support: an IDE element can be marked by a tool as being "dirty" This

CA9-97-003

indicates that the source code for the element must be regenerated because it or its dependencies have changed. The model provides a method to manage this indication shown in Figure 3.

4. associated file operation: the model makes the distinction between contained and associated files. A contained file is simply an IDE element that is structurally related to its container. An example is the source files contained within a project group folder. Contained files are manipulated using the content and structural operations discussed earlier in this section. An associated file is a file reference that is required to further define the IDE element it is associated with. An example of an associated file or header files required in order to use a part, or source files generated from a group or a part. When most cases these are not actually contained within the element with which they are associated.

At the next level of abstraction are the Containment Scope 22 and Name Scope 24 base classes. As mentioned above, containment and part naming are parallel, and therefore, independent relationships.

- 15 On the containment scoping side, the base class 22 represents a generalized grouping construct. In its default form, it allows an arbitrary collection of groups, parts and files that can be owned or linked. Also by default, its content can be "edited" through the tree view or a tool action. This is a general "folder" 40 construct within the IDE.

- 20 The base Name Scope class 24 defines support for nested name scopes within the preferred embodiment model. Its primary intent is scoping of name parts. The Name Scope Folder 26 and Named classes are used to construct the hierarchy reflecting the definitional scopes of the parts. This hierarchy is parallel to that of containment, to allow name scoping to be independent of part containment.

- 25 The Namescope Folder class 26 is in all aspects identical to the Containment Scope class 22, except it also defines a name scope. Namescope Folder 26 is to allow name scope other than name groups and the derivations and parts in their derivations. This is a general names based construct within the

CA9-97-003

IDE.

Scoped names are constructed by assembling a nested set of scopes using the group and part constructs. Model, the singleton basis of the structure, plays the role of the global naming scope.

5 Constructors are provided on group and part derivations that allow specification of the parent name scope for the item being constructed. If a name scope for these objects are not specified the constructed objects are defined at a global scope, that is created below model in the name scoping hierarchy.

10 In general, tree nodes also act as local name scopes for the immediate children they contain. By default, children of an element must have unique names. Since part names are globally unique (under the name scoping mechanism), the default model implementation will allow a part to be created independent of any non-part construct that may exist with the same name but will not allow a part to be stored in an element container if the container already has a non-part child with the same name.

15 Continuing in the Name Scope hierarchy, the Named Part class 28 represents a generalized part construct 44. By default, its content is editable only via an explicit tool action. The part class is not a folder, although it may be defined as an aggregate of sub-parts.

20 The Named Part class models are based constructed type, and permits the data model to capture information about model defined attributes and properties, inheritance, data member and sub-part aggregation, and externally triggerable behaviour or methods. All parts can act as name scopes for other parts.

25 As mentioned above, Primitive Part 30 is a degenerate case of the base Named Part 28. Primitive Part 30 are various primitive types such as integer, character, float, Boolean, supported by the data model.

CA9-97-003

The Builder-Consumable Part class 32 implements semantics of a component part. These builder parts are consumable by the visual composition editor (graphical composer). They are defined in terms of attributes, events and actions. At generation time, Builder-Consumable parts feature maps to the appropriate target component model 46.

5

The Builder-Composed Part class 34 is a partition class that defines the semantics of a partitionable composed part, as distinct from a distributable part. It should be pointed out that a part is distributable if builder tools are allowed to use a proxy to access a remote instance of the part. In the data model, this is indicated by a part property. A part is partitionable if distribution support can effectively rearrange the part content across two or more parts. Concrete parts derived from this class map to composed parts 48 and are partitionable by distribution support.

10

Local associations among a part's sub-parts are captured within the part definition. They represent sub-part interactions. Methods are provided to handle the repartitioning of a part which occurs by redeploing a sub-part from one part to another. If this redeployment is successful, a distributed association between the parts is created. These represent distributed cross part interactions between sub parts of different parts.

15

A final base class is the File Reference class 36 which represents a generalized file reference construct. It is a symbolic pointer to a file item 50 in the file system.

20

Embodiments of and modifications to the described invention that would be obvious to those skilled in the art are intended to be covered by the appended claims.

CA9-97-003

The embodiments of the invention in which an exclusive property or privilege is claimed are defined as follows:

- 1 1. A metadata container for common access tool data in an object oriented programming
2 environment, having a hierarchical structure, comprising:
3 simple constructed types that contain subparts and encapsulated behaviour;
4 components that contain properties of a target language; and
5 composed parts that permit partitioning for distribution.
- 1 2. A metadata store for common access tool data in an object oriented programming
2 environment, comprising:
3 a single base class defining common behaviour for elements in the tool data, and
4 separate abstract class hierarchies, inheriting from the single base class, to define name scopes
5 and containment for the tool data.
- 1 3. A metadata store, according to claim 2, wherein the abstract class hierarchy to define name
2 scopes comprises:
3 a part class representing a generalized part construct adapted to contain subparts and
4 encapsulated behaviour.
- 1 4. A metadata store, according to claim 2 or 3, wherein the abstract class hierarchy to define
2 name scopes comprises:
3 a component class adapted to implement properties of a target language.
- 1 5. A metadata store, according to claim 2 or 3, wherein the abstract class hierarchy to define
2 name scopes comprises:
3 a partition class adapted to contain definitions of part partition semantics

CA9-97-003

- 1 6. A metadata store, according to claim 2, wherein the abstract class hierarchy to define name
2 scopes stratifies the tool data into a layered structure, comprising:
3 simple constructed types that contain subparts and encapsulated behaviour;
4 components that contain properties of a target language; and
5 composed parts that permit partitioning for distribution.

2201278

FIGURE 1




FIGURE 2

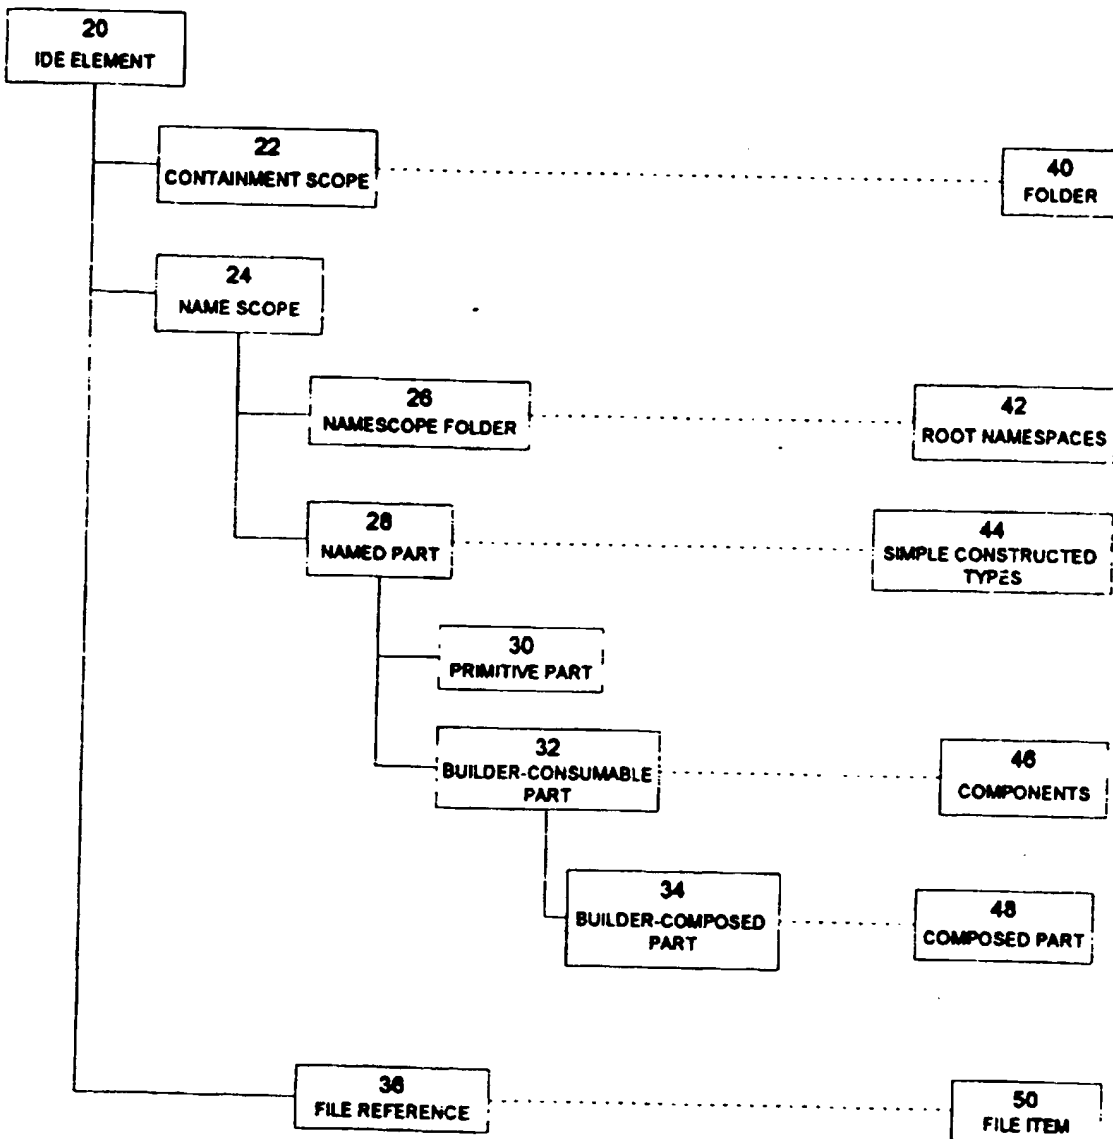
IDE MODEL LAYER HIERARCHYIDE TOOL USAGE

FIGURE 3

```

class DMDEElement public DBase
{
public:
    // Methods to query/manipulate content
    class ContentCursor public DMKeyedOrderedCursor
    {
    public:
        virtual ~ContentCursor () {}
        virtual Boolean isLinked () const = 0
        virtual DMDEElement const& element () const = 0
    };
    DMDEElement ContentCursor* newCursor ();
    Boolean hasContent () const;
    unsigned int numberOfElements () const;
    DMDEElement& elementAt ( const DMDEElement ContentCursor& curs ) const;
    DMDEElement& elementWithKey ( const TString& key ) const;
    DMDEElement* findElementWithKey ( const TString& key ) const;

    // Methods to access/change icon or bitmap
    DMResource getIcon () const;
    void setIcon ( DMResource const& resource );
    DMResource getBitmap () const;
    void setBitmap ( DMResource const& resource );

    // Methods to allow tools to determine if/how content
    // is manipulated
    virtual Boolean canOwn ( DMDEElement& elem ) const = 0;
    virtual Boolean canLink ( DMDEElement& elem ) const = 0;
    virtual Boolean canRemove ( DMDEElement& elem ) const;
    virtual Boolean canMove ( DMDEElement& elem ) const;

    // Methods to manipulate content
    virtual void move ( DMDEElement& fromParent,
        DMDEElement& toParent,
        Boolean own,
        Boolean autoLinking = true );
    void ownAsFirst ( DMDEElement& elem );
    void ownAsNext ( DMDEElement& elem, DMDEElement ContentCursor& curs );
    void linkAsFirst ( DMDEElement& elem );
    void linkAsNext ( DMDEElement& elem, DMDEElement ContentCursor& curs );
    Boolean remove ( DMDEElement& elem );
    void remove ( DMDEElement ContentCursor& curs );
    Boolean remove ( const TString& key );

    // Methods to query or manipulate associated files
    Boolean hasFiles () const;
    void associateFile ( DMFile& file );
    DMFile& fileAt ( const ICursor& curs ) const;
    ICursor* createFileCursor () const;

    // Methods to invoke, enable, disable associated
    // toolactions
    void disableToolAction ( DMToolAction& action ) const;
    void enableToolAction ( DMToolAction& action ) const;
    long perform ( DMToolAction& action, DMSearchCriteria const& context );
    virtual Boolean canEdit () const;
    virtual void edit ( DMSearchCriteria const& context );

    // Method to indicate if tree element is to be shown in the tree
    virtual Boolean isVisible () const;

    // Method to allow and invoke code generation
    virtual void generate ( DMSearchCriteria const& context );
    virtual Boolean canGenerate () const;

    // Methods to query, manipulate the dirtiness state of the
    // tree element object
    void setDirty ( Boolean dirty = true );
    void resetDirty () ( setDirty(false) );
    Boolean isDirty () const;
};

```